arXiv:1606.04287v1 [cs.SE] 14 Jun 2016

Business Matter Experts do Matter: A Model-Driven Approach for Domain Specific Process Design and Monitoring

Adrian Mos, Mario Cortes-Cornax

Xerox Research Center, 6 Chemin de Maupertuis, Meylan, France
adrian.mos@xrce.xerox.com, mario.cortes@xrce.xerox.com

Abstract. Business process design and monitoring are essential elements of Business Process Management (BPM), often relying on Service Oriented Architectures (SOA). However the current BPM approaches and standards have not sufficiently reduced the Business-IT gap. Today's solutions are mostly domain-independent and platform-dependent, which limits the ability of business matter experts to express business intent and enact process change. In contrast, the approach presented in this paper focuses on BPM and SOA environments in a domain-dependent and platform-independent way. We propose to add a domain specific-layer on top of current solutions so business stakeholders can design and understand their processes in a more intuitive way. We rely on previously proposed technical solutions and integrate them in an end-to-end methodology (from design to monitoring and back). The appropriateness and the feasibility of the approach is justified through a use case and a complete prototype implementation.

Keywords: Model-driven methodology, process monitoring, DSL, BPM, SOA

1 Introduction

Business process design connected to execution and monitoring are critical for successful Business Process Management (BPM) [21]. Today, the Business Process Model and Notation [16] (BPMN 2.0) has become the de-facto standard for business process modelling. With the aim at filling the Business-IT gap, significant effort has been put into bringing BPMN executable and closer to Service Oriented Architectures (SOA). A BPM Suite (BPMS) manages the process execution directing SOA calls to the appropriate services and generally provides monitoring infrastructure. While these components help alleviate agility problems that business stakeholders encounter, there are important limitations to the current approaches. We observed that most of the existing solutions are domain-independent and platform-dependent, which limit the power of business matter experts at the design and monitoring stages.

Concerning **process design limitations**, the BPMN standard lacks guidance to reach executable processes from high-level process models. Silver [24]

highlights this problem, and proposes a level-based top-down approach to design business processes (*Descriptive level*, *Analytical level* and *Execution level*). However, the generality of the most common BPMN 2.0 graphical elements, in particular the Task element, reduces semantic expressiveness [12]. Business analysts require dedicated means (e.g., specific type of task with implicit domain knowledge) to effectively model their business domain (ex. logistics, healthcare, transportation, etc.) [18]. Domain Specific Languages (DSLs) are an effective means to deal with these problems, providing improvements in expressiveness and ease of use [11]. More specifically, Domain Specific Process Modelling Languages (DSPMLs) [6] permits business stakeholders to design their processes in a much more intuitive way than BPMN.

Regarding **monitoring limitations**, BPMS solutions collect and present data at the level of the process description, which is generic. This fact results in monitoring information that is collected in a generic way with respect to the business domain (ex., "activity", "gateway" or "event") with no correlation with the business concepts (ex."order handling" or "shipping") apart from the simple matching "label - BPMN element". This causes a number of problems: (1) it is hard to make use of the monitoring data in order to present meaningful metrics for business users, without significant configuration efforts for each BP; (2) it is difficult to correlate the business concepts to the execution of services in the SOA layer; (3) it is difficult to set wide-ranging SLAs that affect all BPs in the organization equally. For instance, it may be necessary to specify that all the "shipping" operations, regardless the BP in which they occur, must execute in less than 2 days.

In this paper, we present an approach that focuses on BPM and SOA environments in a domain-dependent and platform-independent way. Previous technical solutions [13–15] are combined to present a methodological, model-driven approach that integrates domain specific modelling with domain specific monitoring in an end-to-end solution. The appropriateness and the feasibility of our approach is shown through a use case and a complete prototype implementation. The rest of the paper is structured as follows. Section 2 describes a general overview of our method, based on a running example. Section 3 details the steps of the approach. Section 4 focuses on the prototype implementation. Section 5 presents related work and finally, Section 6 concludes and discusses future work.

2 Overview of the Approach

Figure 1 gives an overview of the approach from a modeller (business analyst and architect) point of view. Each number corresponds to one key step that will be further described in the following sections. The figure contains a simplified order handling process. The orders are received either by a submission web form or by standard mail. In the latter case, some document pre-processing is necessary in order to handle the order (i.e., scanning, Optical Character Recognition (OCR), and segmentation to extract the different sections). The order's comments, which could be in different languages, need to be handled before the approval. After-

wards, some classical processing steps such as payment, packaging, preparation of the documents (i.e., tracking number, bill), as well as the actual shipping and the confirmation are defined. In dotted lines, the business stakeholder indicates the exceptional paths. Each symbol represents a *Domain Concept*, which makes reference to an enterprise well defined know-how element.

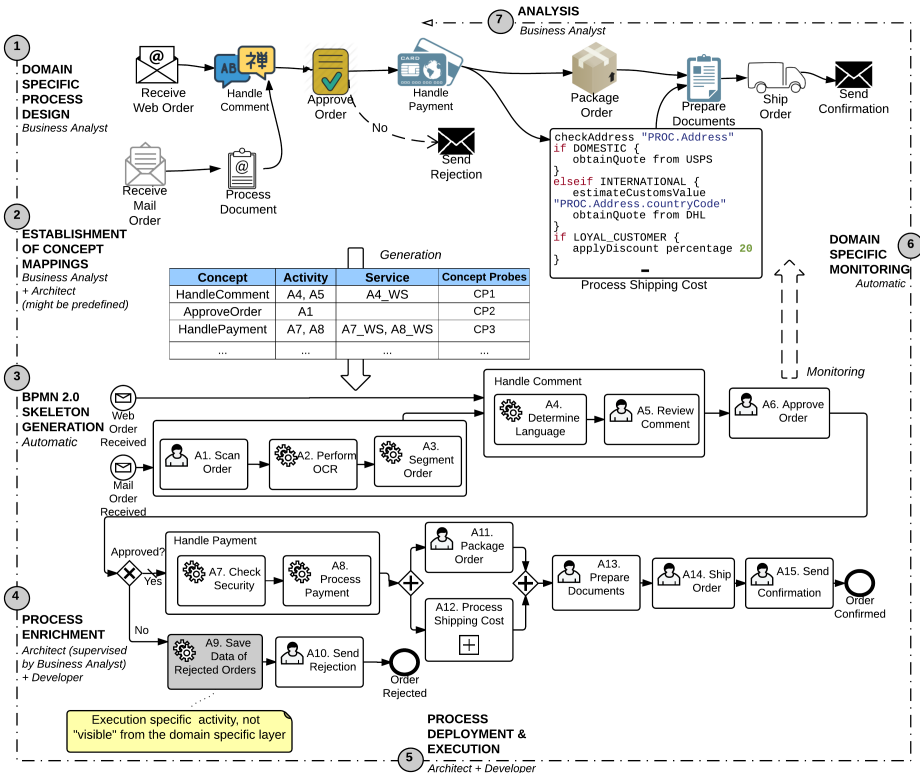

Fig. 1. Approach Overview with Main Steps

The first step corresponds to the **domain specific design**, using a DSPML (*Step 1* in Fig. 1). The domain specific language must have been previously designed based on the generic domain meta-model that we propose. Potentially, several DSPMLs can be combined, as the example shows. For instance, in order to define the process of calculating the shipping cost, a textual description may be more appropriate. In the graphical part, we advocate taking into account the principles of notations defined by Moody [12]. More details about a particular language are out of the scope of this paper. However, rich language definitions are possible for various domains, as we show also in [14].

The analyst can then **establish the concept mappings** (*Step 2* in Fig. 1). While business concepts are already connected by default to the abstract services from the enterprise repository, the links can still still be modified at this stage.

This is essential in grounding the domain knowledge in technical realities. For instance, "Handle Payment" corresponds to two technical services. It will imply the creation of the corresponding service tasks in BPMN. The mapping between domain concepts with the process activities relies on a pivot meta-model and unique ID (UID) attributes. The so-called Common meta-model (CommonMM) is a central, simplified representation of the main generic process concepts common to business process descriptions, such as activity, flow and gateway. It is significantly simpler than fully-fledged BPMN because its objective is simply to extract the essence of the structure of various business processes. Our hypothesis is that a descriptive level [24] (reduced amount of symbols but semantically enriched) is enough to define high-level domain-specific process models.

The **BPMN 2.0 skeleton is generated** relying on the aforementioned CommonMM (*Step 3* in Fig. 1) and the concept mappings (see table with concepts mapped to activities). Note that the concept-to-activity mappings are generated or validated at this stage. Also note that the transformations are transparent to the business analyst. At most, the latter will have to agree with the business architect on the correspondence between the domain concepts and the to-be activities supported by generation templates. Once transformed into an instance of the CommonMM, the processes can be converted to the process modelling language of choice.

Generated BPMN models are typically enriched and refined (*Step 4* in Fig. 1). Extra activities, a complete data model or a resource model may be necessary in order to enable executability.

Deployment and execution follow (*Step 5* in Fig. 1). The only constraint that we impose here is the preservation of the concept mappings (i.e., not manually deleting the generated UIDs). Extra activities that may be added in the BPMN are considered as technical additions and of reduced interest from a business point of view (ex. activity A9). These activities will not be represented at the DSPML level when showing information coming from the domain specific monitoring infrastructure. The deployment phase is necessary to install the process artefacts and bind the generated abstract services with actual services, which will be running for instance in an enterprise service bus.

Monitoring (*Step 6* in Fig. 1) aims at aggregating and displaying data in the domain specific environment relying on information from the concept mappings. Our proposition aims to address the aforementioned shortcomings of today's monitoring capabilities for BPMS/SOA applications. A layer of abstraction is added on top of the existing capabilities rather than replacing them. The platform-independence ensures compatibility with a wide range of existing systems and platforms.

Finally, in an **Analysis** stage (*Step 7* in Fig. 1) the monitored data is studied to iteratively improve the process. Iteration may also imply the enrichment of the enterprise know-how, which is capitalised through the domain concepts.

To summarise, the interest of the contribution is twofold: 1) the approach takes into account in a very specific way the business stakeholders, enabling domain specific modelling and monitoring; and 2) the entire cycle is integrated

in a continuous improvement approach, supported by tools through model-driven transformations.

3 A Model-driven Approach for Domain Specific Process Design and Monitoring

This section details the main ideas of our model-driven approach for domain specific process design and monitoring, which considers business stakeholders as first class-citizens for BPM. The section mainly focuses on *domain-specific design*, *establishment of concept mappings* and *domain-specific monitoring* which are the most relevant part of the work. The *BPMN generation*, the *process enrichment*, the *deployment and execution* and the *analysis*, while implemented and integrated, are not described in much detail, as they are relatively common BPM activities.

3.1 Domain-Specific Design through Domain Concepts

The interest and the limits of DSPMLs have already been presented in previous work [13]. Naturally, our goal is not to propose a particular DSPML as their aim is to be adapted to particular business needs. Instead, we propose a generic domain description meta-model (MM), which provides a structural view of the domain. We then illustrate it with examples corresponding to the use-case described in the previous section.

The upper part of Fig. 2 provides in a simplified way, the meta-models used to define the key points of the business domain in a generic way. They represent business domain information for an enterprise, with regard to the specification of concepts that are going to be reused in the business processes. The domain meta-model is useful for several proposes: 1) to store the domain information in a central repository on a collaboration and distribution server. This allows common access to the defined concepts to all the business users; 2) to generate a domain editor (textual) that can be used stand-alone or embedded in a graphical editor as part of a diagram designer; 3) to make the connection with the pivot meta-model specifying how process steps are going to be represented. This point is important when in a diagram, the user specifies that a process step is going to perform a business function corresponding to a business concept; 4) to inform and update SLA for business concepts. An enterprise-wide SLA management ensures that all activities and all processes that refer to a particular business concept would be marked with appropriate SLA constraints. This can bring important advantages when changes to company policies have sweeping implications on many SLAs, as they can be automatically propagated to all the relevant activities and processes.

The meta-model in Fig. 2 defines a *Domain*, which contains a set of domain specific concepts (*DSConcept*). A *Domain* also contains *SLA* elements, describing the agreement details. A *DSConcept* relates to *DSService* elements describing the actual SOA services required in the domain. Note that services can be abstract entities bound later in the deployment phase [7].

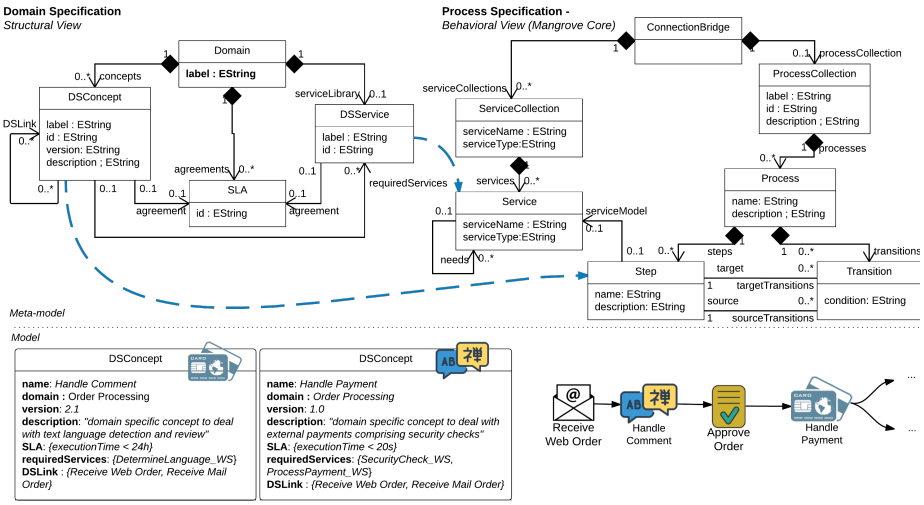

Fig. 2. Domain Specific Concepts Design Relies on a Generic Domain Meta-model

Illustration based on the use case. A Domain Concept supports the representation of business domain knowledge in an enterprise. Fig. 2 illustrates in the bottom part how knowledge common to the enterprise is stored in two example domain concepts (in contrast to a pure BPMN approach where such information would be implicit in the minds of the designers). These concepts would typically be stored in shared repositories. The information comprises for instance a name (verb+object) a version number and the SLA. Links between domains concepts are defined in order to define dependencies. For example, the concepts: "Handle Payment" and "Handle Comment" are related to the "Receive Mail Order" and the "Receive Web Order" business activities. This means that a common DataObject will be shared between the BPMN activities that are generated. Note that the data-model generation is currently not supported by our solution, although it is being investigated. However, these links provide necessary hints to the architects and analysts that enrich the generated BPMN skeleton.

3.2 Establishment of Concept Mappings

In their simplest form, concept mappings are connections between business concepts and the SOA services that are used by them. This relation, could be defined by means of process activities. A simple example of the concept mappings is presented in Fig. 2. Concept mappings are defined as following:

- Set of services $S = \{s_1, s_2, ..., s_q\}$
- Set of processes $P = \{p_1, p_2, ..., p_m\}$
- For each process p_k, a set of activities $Ak = \{a_{k1}, a_{k2}, ..., a_{kt(k)}\}$ where the number of the activities in the set t(k) depends on the complexity of p_k
- The set of all activities in all processes $A = A_1 \bigcup A_2 \bigcup ... \bigcup A_{|P|}$

The goal of concept mappings is to determine the following sets:

- Set of concepts $C = \{c_1, c_2, ..., c_n\}$
- $ConceptMappings(CM) = \{c_j, s_j : \forall c_j \in C; S_j \subseteq S\}$ which contains for each concept its list of services, e.g., $HandlePayment, (s_1, s_2)$.
- $ActivityMappings(AM_k) = \{a_k, c_j : \forall a_{ki} \in A_k; c_j \in C\}$ which contains for process p_k its activities and the concepts they map to.
- $AM = AM_1 \bigcup AM_2 \bigcup ... \bigcup AM_{|P|}$ which contains for each activity all processes the concept it maps to.

Obtaining the sets C and CM requires that the business concepts used in the processes be clearly identified together with their required SOA services. Concepts are defined by business experts, connected to abstract services initially and eventually bound to real SOA services in the deployment stage as discussed in Section 3.5. The modelling environment needs to propose to the business expert a set of relevant SOA services. Other approaches, more or less automatic, for concept mapping could be applied [13]. Once the concepts have been identified, it is necessary to obtain the AM set by mapping the BP's activities to the concepts (typically done automatically at the BPMN generation phase).

Illustration based on the use case. A concept can have an immediate correspondence with a process activity (ex. "Approve Order") or several activities (ex. "Handle Comment", which refers to *Determine Language* and *Review Comment* in Fig. 1). A domain concept which is described with a textual DSL as for example the "Process Shipping Cost", corresponds to a sub-process that will generate several BPMN activities. The sub-process itself contains a number of domain concepts that correspond to the knowledge about price management. These correspondences will vary depending on the enterprise domain concepts. In fact, the freedom to define such mappings brings an important level of flexibility in how business knowledge gets transferred into processes that are governed in a uniform way at the business level.

3.3 BPMN Skeleton Generation

The BPMN generation relies on a Common Meta-model (CommonMM), which is a simplified representation of the main generic process concepts. The reason why the DSPML-based processes are not directly transformed into the generic language is to introduce flexibility in the approach. As the generic language (usually BPMN) evolves, only the transformation between the CommonMM and the target MM needs to be updated. It could be argued that the use of a simplified version of the BPMN meta-model, where only the descriptive objects are included, could facilitate the transformation process. However, if we aim to strictly follow the BPMN 2.0 meta-model, a complex class hierarchy should be respected. This particularity may not be shared with other process languages and would complicate transformations (ex. a *Task* element subsequently inherits from *Activity*, *Flow Node*, *Flow Element*, and *Base Element*). For our prototype, we use

Mangrove Core[1] as our CommonMM (a simplified version is depicted in the upper right side in Fig.2). Mangrove Core is a meta-model that unifies business processes and SOA elements. It provides behavioural support to the domain definition in order to define the necessary steps in a process. This framework, does not aim to manage a large collection of processes, such as APROMORE [20]. Instead, it focuses on preserving the sync between the common elements of business processes and architectural constructs from the various related diagrams.

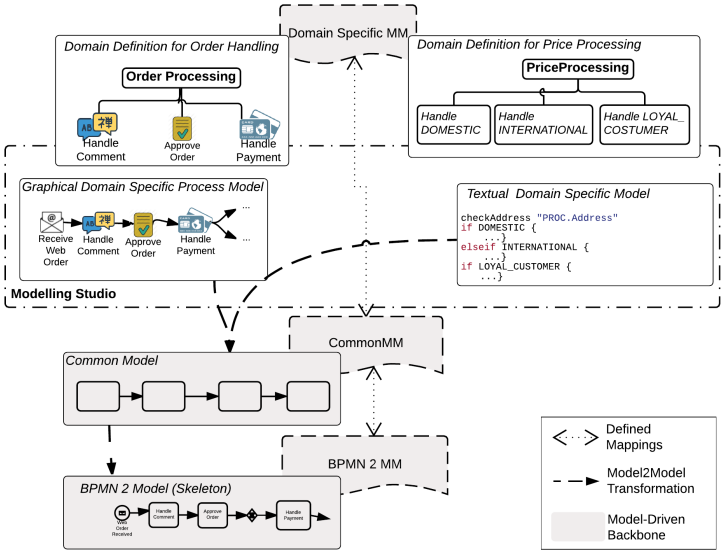

Fig. 3. Model Transformation from Domain Specific Models to BPMN 2.0 Model

In our approach, several target languages can be supported incrementally over time. When a new target is added, new transformations need of course to be added between the CommonMM and the new MM. We do not go into the details about the transformation process as the paper focuses on this general methodology. More details about the two-way synchronization between domain specific models and BPMN are presented in [14].

Figure 3 depicts how the BPMN generation is performed through model transformations. In our running example, two domain specific models are merged in a unified BPMN model. This shows the capacity of adaptation of the approach to different business expressibility needs. Both domain specific models leverage the DomainMM, which is mapped (i.e., MM concepts are linked) to the CommonMM. The latter is mapped to the target meta-model (in this case BPMN). The depicted meta-models provide a model driven backbone, where different domain specific models can be plugged in. The modelling studio (see Fig. 3) is the tool that permits the business analysts to build specific process models connecting the predefined domain concepts using various process representations, based on their specific business domain.

[1] http://www.eclipse.org/proposals/mangrove/

3.4 Process Enrichment

This stage relates to the need of the generated BPMN skeleton to be enriched if execution is targeted. New activities (ex. A9 in our running example), specific gateways and events, as well as several details may need to be added to the process model. We do not go into much details here as our approach does not impose any significant restrictions to this stage. The only constraint that the approach brings is to preserve the generated activities (tasks or sub-processes), so the link between the domain concepts and the process activities be maintained. Indeed, we did not force a perfect vertical alignment that could be very costly and unrealistic as described in [25]. The double synchronization mechanism explained in previous work [14] permits to make (and propagate) changes in the domain model as well as in the generated BPMN model. The tracking of generated elements can be based on several identity-preserving mechanisms, of which a simple example is the usage of unique IDs injected in hidden properties of BPMN elements. This mechanism enables the possibility to make changes in the domain model as well as in the generated model.

3.5 Process Deployment and Execution

Concerning deployment, when defining business processes, individual business process activities can be connected to the service-execution capabilities of the enterprise, thus allowing any business process to be easily translated into an executable workflow on the platform of choice. This capability is enabled in our approach by providing mappings for each domain concept in order to specify how it should be grounded in the SOA. These mappings are done with idealized or abstract services in a two-step mechanism, in order to ensure better portability (and reusability) across the enterprise, as well as encourage proper adoption of good SOA-practices in future evolutions of the enterprise SOA. These abstract services (AS) would then be further connected to the real services in the repositories. The creation of these mappings would typically be performed by IT experts that have a good understanding of the domain and who envisage an ideal connection to a SOA. These abstract, idealised services, would not necessarily correspond on a 1 to 1 basis with business concepts as we show in the example. That is because there are important differences in concerns when defining business elements and when defining the service infrastructure, due to varying needs for reusability, performance and evolution of these two layers. In our approach, this two-step binding mechanism explained in [15] is applied to link domain concepts to any number of AS first and then each AS to real SOA services.

In order to execute the process, BPMS usually need at least a data-model defining the artefacts that flow in the process, a resource-model establishing the links between the roles defined in the process and actual users and the implementation of gateway conditions (usually based on data). These artefacts can be partly generated by the presented approach, but they may need to be enhanced by technical architects.

3.6 Domain-Specific Monitoring

The main elements involved in domain specific monitoring are the Concept Probes (CPs) and the Business Process Probes (BPPs). There is a one-to-one correspondence between CPs and domain concepts. CPs collect an arbitrary number of metrics, such as execution time or execution status from the activities that are mapped to a domain concept. Once the CPs are created, they need to be bound to the monitoring capabilities of the existing infrastructure, effectively acting as an extra monitoring layer on top of the actual BPMS and SOA platforms. BPPs aggregate data from the BPMS and the various CPs. In order to enable them, they have to be linked to the domain concepts at design time (which is performed automatically). When all the required mappings are available, the probes are created, instantiated and deployed automatically respecting a predefined template. More technical details about concept probes can be found in [13]. Here, we summarise their main functionality. Both CPs and BPPs are divided in three main components with particular concerns: the *Raw Data Collection Component*, mainly collects data from the activities corresponding to each concept and the related technical services. The *Analysis Component* is in charge of the aggregation of a raw data into composite metrics. These composite metrics are data structures that present the aggregate monitoring information combining the individual metric data for BPMS, SOA and other collection points such as Network Monitoring, App Server Monitoring and Operating System Monitoring. Finally, the *Alerts Reporting Component* allows the registration of SLA requests through a configurable alert port. It uses the analysis component to constantly compare the aggregated metric values with the required thresholds.

The approach provides the business stakeholders with means to govern their processes at a high level, with impact to the entire collection of business processes in a domain, if required. Relying on domain concepts, they are able to consistently manage the execution parameters of a large collection of process descriptions and their instances. For example, if the Shipping concept is already defined, it is automatically reused in any process description detailing shipping operations, carrying over the reuse of the generation and the monitoring infrastructure. The definition of corporate-level SLA is easily implemented and maintained. Relying on the generative approach, changes are spread through the different layers. In the long run, the monitoring mechanisms enable better decision making, based on domain specific information, by putting the appropriate level of information in the tools used by the business-matter experts. Section 4 discusses the prototype implementation and provides more details on the actual set up of the monitoring probes.

3.7 Analysis

In order to close the iterative lifecycle loop depicted in Fig. 1, an analysis step is necessary (*Step 7* in the figure), where the analysts study the monitored data in order to improve the process. The novelty in our approach is that the new know-how acquired in the enactment of the process may imply the update or

creation of domain concepts. One of the biggest advantages of the approach is that if an updated concept is being used in a collection of processes, the changes will more easily propagated through the complete stack.

4 Prototype Implementation and First Validation Steps

Fig. 4 depicts the architecture of the prototype illustrated for our use case. The picture shows the domain specific layer as an additional layer to the BPM and SOA stack. A domain specific editor would be the entry-point for a business stakeholder, providing domain specific process design (based on domain concepts), BPMN generation (which is transparent to business stakeholders) and display of monitoring result (outcome of the concept and process probes). We present some key points of the prototype implementation supporting the process life-cycle. This prototype is mostly based on Eclipse technologies, which are highly relevant in the BPM landscape as many BPM suites are actually built using the Eclipse platform. The discussion relies on the seven steps of our model-driven methodology.

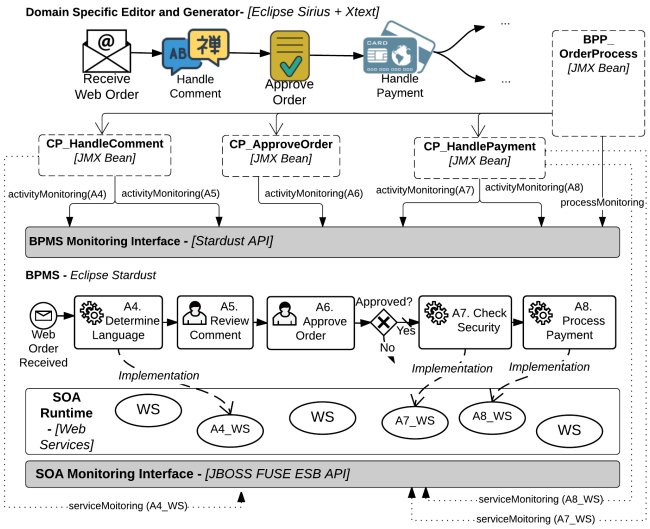

Fig. 4. Prototype Architecture

Process design and concept mappings. The Eclipse Modelling Framework[2] is used for the definition of the domain-specific meta-models. Ecore meta-models are the inputs for the Sirius toolkit[3], which allows rapid creation of graphical domain-specific modelling studios. Fig. 5 shows a screenshot of the graphical studio. It depicts how concepts from the domain palettes can be used to compose processes that have predefined SOA connections to domain services.

[2] http://www.eclipse.org/modeling/emf/
[3] https://eclipse.org/sirius/

Monitoring information can be shown in various ways, in this particular example, execution times in the process elements indicate the BPMN activities' contribution to the overall execution time. The service contribution time is indicated in the DSConcept-Service links (dotted lines). Today, the creation of the domain-specific editor has to be supported by technical architects and developers. However, we are working on a generative approach that permits to dynamically create these modelling editors from the definition of the domain concepts.

BPMN 2.0 generation. In addition to the Mangrove Core meta-model that we use, the Mangrove project provides a variety of plugins for model transformations as well as some editor extensions. The model-transformation plug-ins contain code that convert supported meta-models to Mangrove Core and vice-versa. They are invoked from editor plug-ins that are connected to the supported editors through standard extension points. Note that the generator only outputs the model definition and not the visual layout of the model. The BPMN 2 Modeller[4] is used to initialise the graphical representation from the generated model with a built-in Mangrove support wizard.

Enrichment of BPMN models, deployment and execution. In our scenario, the generated BPMN skeleton is further enriched with a simple data-model, the implementation of the gateway conditions and a resource-model in order to enable execution. As Fig. 4 indicates, we use the Eclipse Stardust[5] BPMS to execute our process. The choice of this BPMS was made because of the maturity of the tool and openness of its process monitoring API, which easily allows access to detailed process monitoring information from external components (our concept probes).

Monitoring. The generation of the concept probes is done through template instantiation. Once they are generated they need to be managed as components managed by the monitoring framework. We use the Java Management Extensions (JMX)[6] for our distributed monitoring infrastructure managing the probes as well as for integrating with existing monitoring frameworks. JMX is supported by a large variety of infrastructures, both commercial and open-source.

As an initial validation step, we rely on the SEQUAL (SEmiotic QUALity Framework) [8], which is widely used and goes beyond the modelling language to characterise its quality. We conclude that the proposed approach can significantly complement other BPMN approaches regarding the SEQUAL framework: the *domain*, *comprehensibility*, and *organisational* appropriateness are improved by the fact that the actual focus is specific to the domain. Indeed, the framework advocates that a language must be powerful enough to express anything in the domain but no more. Also, a language should be easily extensible in order to adapt to changing business needs. These points clearly justify the interest of a DSPML on top of a BPMN model. The *modeller* appropriateness and the *participant* appropriateness will not change significantly as we propose to ultimately

[4] https://www.eclipse.org/bpmn2-modeler/
[5] https://www.eclipse.org/stardust/
[6] https://docs.oracle.com/javase/tutorial/jmx/

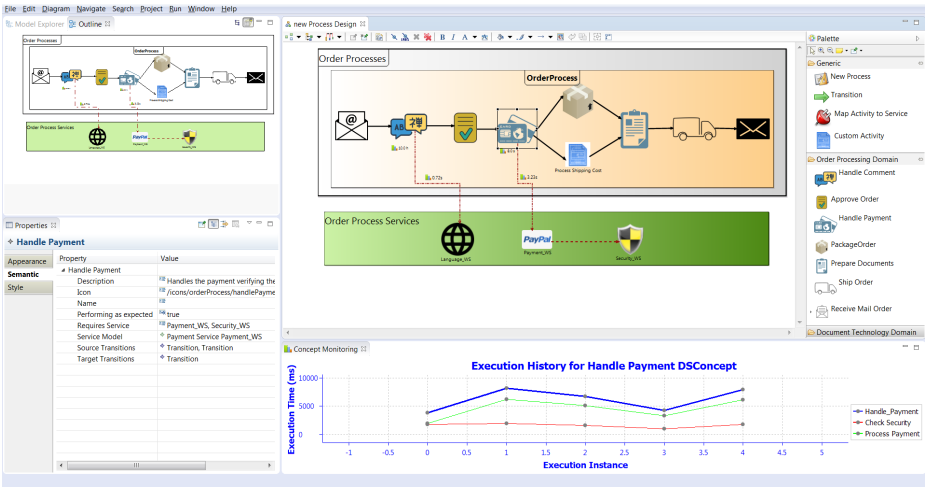

Fig. 5. Screenshot of the Eclipse-based graphical studio

rely on BPMN. In fact, the framework recommends the use of well-known modelling languages and our approach targets basic BPMN generation. Finally, the use of proven model-driven technologies such as Sirius permit a good *tool* appropriateness. Obviously, these improvements will highly depend on the proposed DSPML, but the approach provides the means to achieve them. Qualitative evaluations with final users are envisaged in order to complete the validation. Practical experiments may result in changes or refinements of the approach. The method could be extended to incorporate an user-centred approach to build the DSPML as discussed in [19].

5 Related Work

Related work can be analysed from two main aspects: the model driven approach and the monitoring capability. Related to the **model-driven** part, Heitkotter [4] proposes DSLs4BPM, an approach for creating domain-specific process modelling languages. On the same line, Grundy et al. [3] rely on Eclipse tooling to propose domain specific visual language editors. The difference with our approach is that these works do not provide a structured methodology to design and analyse the processes as we do. More important, the monitoring part, which is essential for business experts, is not considered. Becker et al. [2] propose the modelling method called PICTURE, which specially focuses on public administration processes. The so-called "process building blocks" could be compared with our domain concepts, as they are high-level domain-specific artefacts that help build the actual process. Kumaran et al. [9] follow the same line, proposing to automate complex and variable workflows in a service delivery management architecture. The main difference is that our approach can leverage BPMN solutions (the de-facto standard) in order to reach execution and monitoring. Other works propose extensions to BPMN 2.0 in order to be domain-specific [22]. These

approaches are limited by their focus on a very concrete problem space while still having to deal with the aforementioned complexity and generality of BPMN 2.0. Goal-oriented approaches [10, 23] use goal models as a preliminary step for process modelling. However, the graphical notations of the more popular goal oriented languages (i*, KAOS and MAP) still lack of *Semantic Transparency* [12]. There is also limited tool support for goal modelling. In addition, the goal-driven generation approaches tend to propose goal models closely tied to the business process.

Considering the **monitoring** part, there are approaches that recur to aggregation mainly to compose events from a low-level monitoring source (using Complex Event Processing queries) in order to extract more meaningful data out of raw events [5, 17]. Such approaches use a variety of techniques to derive better understanding of raw events, but they fundamentally still stay at a generic level with regard to the business domain. There are also approaches that try to correlate execution events to the originating processes using some forms of traceability between model elements and execution events. For instance, in [1], the authors argue for the existence of domain-specific patterns for interpreting events, without giving a complete solution. Their suggestion is in line with our proposition in the idea of presenting information corresponding to domain elements, but they mostly focus on interpreting CEP events, while our approach targets structured probes that connect directly with monitoring APIs. In summary, the studied approaches recur to generic event analysis and do not provide a "native" monitoring probe layer that directly correspond to the business concepts. To the best of our knowledge, there is no work providing an end-to-end solution for domain specific process design and monitoring.

6 Conclusion and Future Work

Existing design and monitoring approaches are typically technology-specific and generic with respect to the business domain. This limits the ability of business matter experts to express their intent and enact process change. This paper leverages current BPM and SOA solutions adding a layer that is domain-dependent and platform-independent in order to facilitate process design by business matter experts. The approach presented in this paper also simplifies the management of complex business processes that span multiple domains of expertise through the support of several domain definitions during process design.

We have presented a methodological and iterative approach that relies on seven main steps : 1) *domain specific design* using the so-called domain concepts, which comprise the explicit representations of enterprise domain know-how; 2) *the establishment of concept mappings* between domain concepts and process activities and technical services; 3) *BPMN generation* relying on a pivot meta-model that enables flexibility and facilitates model transformations; 4) *process enrichment*, which does not seek perfect vertical alignment between high-level models and executable ones but keeps artefacts in sync relying on concept-mappings; 5) *deployment and execution*, which defines a two-step bind-

ing mechanism between domain-concepts and technical services; 6) *domain specific monitoring*, based on Concept Probes and Business Process Probes that map service and process monitoring metrics to the domain concepts and 7) the *analysis* of the monitored data, which may imply the enrichment of the domain concept repository. The presented methodology is supported by tools that automate the generation and synchronisation activities. We used a mature set of open-source tools from the Eclipse Ecosystem to implement a fully functional prototype and used a running example throughout the paper to illustrate the interest and applicability of our proposition.

We are focusing our next explorations on the following three main points. Firstly, the automatic generation of graphical process model editors from domain specifications mapped to the definition of the abstract syntax of the language and additional functional templates. Secondly, the integration of collaborative modelling in the aforementioned editors, which is critical in business process design. Thirdly, the automatic generation of various artefacts for the process data-model that could be used in the actual process implementation. We also aim to connect the data-model to the monitoring probes in order to correlate execution information to process data flow. These points are all under advanced stages of exploration, with a prototype being developed using Eclipse-based open-source technologies.

References

1. Ammon, R.v., Silberbauer, C., Wolff, C.: Domain specific reference models for event patterns–for faster developing of business activity monitoring applications. In: VIP Symposia on Internet related research with elements of M+ I+ T+. vol. 16 (2007)
2. Becker, J., Pfeiffer, D., Räckers, M.: PICTURE-a new approach for domain-specific process modelling. In: CAiSE Forum. pp. 11–15 (2007)
3. Grundy, J., Hosking, J., Zhu, N., Liu, N.: Generating domain-specific visual language editors from high-level tool specifications. In: 21st IEEE/ACM International Conference on Automated Software Engineering (ASE). pp. 25–36. IEEE (2006)
4. Heitkötter, H.: A framework for creating domain-specific process modeling languages. In: 4th International Conference on Software and Data Technologies (ICSOFT). pp. 127–136 (2012)
5. Hummer, W., Inzinger, C., Leitner, P., Satzger, B., Dustdar, S.: Deriving a unified fault taxonomy for event-based systems. In: Proceedings of the 6th ACM International Conference on Distributed Event-Based Systems. pp. 167–178. ACM (2012)
6. Jablonski, S., Volz, B., Dornstauder, S.: Evolution of business process models and languages. In: 2nd International Conference on Business Process and Services Computing (BPSC). pp. 46–59. Citeseer (2009)
7. Jacquin, T., Mos, A.: Deployment of business processes in service-oriented architecture environments (Apr 28 2015), US Patent 9,021,420
8. Krogstie, J., Sindre, G., Jørgensen, H.: Process models representing knowledge for action: a revised quality framework. European Journal of Information Systems 15(1), 91–102 (2006)

9. Kumaran, S., Bishop, P., Chao, T., Dhoolia, P., Jain, P., Jaluka, R., Ludwig, H., Moyer, A., Nigam, A.: Using a model-driven transformational approach and service-oriented architecture for service delivery management. IBM Systems Journal 46(3), 513–529 (2007)
10. Lapouchnian, A., Yu, Y., Mylopoulos, J.: Requirements-driven design and configuration management of business processes. In: 5th International Conference on Business Process Management (BPM), pp. 246–261. Springer (2007)
11. Mernik, M., Heering, J., Sloane, A.M.: When and how to develop domain-specific languages. ACM computing surveys (CSUR) 37(4), 316–344 (2005)
12. Moody, D.L.: The "physics" of notations: toward a scientific basis for constructing visual notations in software engineering. IEEE Transactions on Software Engineering 35(6), 756–779 (2009)
13. Mos, A.: Domain specific monitoring of business processes using concept probes. In: Service-Oriented Computing – ICSOC Workshops. pp. 213–224. Springer (2015)
14. Mos, A., Jacquin, T.: Improving process robustness through domain-specific model transformations. In: 17th International Enterprise Distributed Object Computing Conference Workshops (EDOCW). pp. 188–193. IEEE (2013)
15. Mos, A., Jacquin, T.: A platform-independent mechanism for deployment of business processes using abstract services. In: 17th IEEE International Enterprise Distributed Object Computing Conference Workshops (EDOCW). pp. 71–78. IEEE (2013)
16. OMG: Business process model and notation (BPMN) version 2.0 (2011). Available on: http://www. omg. org/spec/BPMN/2.0 (2011)
17. Pedrinaci, C., Lambert, D., Wetzstein, B., Van Lessen, T., Cekov, L., Dimitrov, M.: Sentinel: a semantic business process monitoring tool. In: Proceedings of the first international workshop on Ontology-supported business intelligence. p. 1. ACM (2008)
18. Pinggera, J., Zugal, S., Weber, B., Fahland, D., Weidlich, M., Mendling, J., Reijers, H.A.: How the structuring of domain knowledge helps casual process modelers. In: Conceptual Modeling (ER), 2010, pp. 445–451. Springer (2010)
19. Rieu, D., Santorum, M., Movahedian, F., et al.: A participative end-user method for multi-perspective business process elicitation and improvement. Software & Systems Modeling pp. 1–24 (2015)
20. de la Rosa, M., Reijers, H.A., Van Der Aalst, W.M., Dijkman, R.M., Mendling, J., Dumas, M., García-Bañuelos, L.: Apromore: An advanced process model repository. Expert Systems with Applications 38(6), 7029–7040 (2011)
21. Rosemann, M., vom Brocke, J.: The six core elements of business process management. In: Handbook on Business Process Management 1, pp. 105–122. Springer (2015)
22. Saeedi, K., Zhao, L., Sampaio, P.R.F.: Extending BPMN for supporting customer-facing service quality requirements. In: 17th IEEE International Conference on Web Services (ICWS). pp. 616–623. IEEE (2010)
23. Santos, E., Castro, J., Sanchez, J., Pastor, O.: A goal-oriented approach for variability in BPMN. In: Workshop em Engenharia de Requisitos (WER) (2010)
24. Silver, B.: BPMN Method and Style: A levels-based methodology for BPM process modeling and improvement using BPMN 2.0. Cody-Cassidy Press, US (2009)
25. Weidlich, M., Barros, A., Mendling, J., Weske, M.: Vertical alignment of process models – how can we get there? In: Enterprise, Business-Process and Information Systems Modeling: 10th International Workshop (BPMDS). pp. 71–84. Springer (2009)